

HUMAN FACTORS LESSONS LEARNED FROM A SURFACE MANAGEMENT SYSTEM SIMULATION

Deborah Walton*, Cheryl Quinn[†], Stephen Atkins[‡]
NASA Ames Research Center
Moffett Field, CA 94035

ABSTRACT

The Surface Management System (SMS), being developed at NASA Ames Research Center in conjunction with the FAA, is a decision support tool that helps air traffic controllers and air carriers manage aircraft movements on the surfaces of busy airports. By presenting information and advisories to the Air Traffic Control Tower, Terminal Radar Approach Control (TRACON), En-route Center, and air carriers, SMS creates shared departure situational awareness, thereby increasing the efficiency, capacity, and safety on the airport surface. This paper discusses the human factors lessons that were learned during the real-time simulation of SMS that was conducted in January, 2002 at the Future Flight Central facility at NASA Ames Research Center. Five active Dallas-Fort Worth (DFW) Tower controllers participated in the simulation. This paper discusses one of the main objectives of the three-day simulation; to conduct human factors studies to better understand the challenges involved with introducing automation tools into the Tower environment. To this end, studies observing controller coordination were conducted, usability, suitability, and acceptability questionnaires were administered, and informal debriefs were held after each of the nine runs.

BACKGROUND

Tower controllers are responsible for the safe, orderly, and expeditious flow of air traffic on the airport surfaces. Specifically, Tower controllers currently are responsible for taxiing aircraft, sequencing aircraft for departure, and clearing flights for takeoff and landing. In order to get information about the current state of the aircraft and airport resources, Tower controllers use several different information sources: flight strips, a map display known as the Airport Surface Detection Equipment (ASDE), and a repeater of the Terminal Radar Approach Control (TRACON) radar, known as a Digital Brite Radar Indicator (D-BRITE).

Flight strips provide detailed flight information for each departure aircraft including the aircraft type, first departure fix, flight plan, and flight ID of the aircraft.

The D-BRITE provides controllers with the flight identification (ID) of aircraft in the terminal airspace. The ASDE, which presently operates at many large airports, provides a map display of the airport surface that shows the locations of aircraft and other vehicles. The map display provides aircraft location information in an intuitive display that is similar to the controllers' out-the-window view.

ASDE does not identify the aircraft flight number or provide any other flight-specific information because it is a primary (i.e., skin paint) surface surveillance radar. New surface surveillance systems, such as ASDE-X and a prototype that is being developed under the FAA's SafeFlight 21 program, will provide real-time information about the location and identity of aircraft. The ASDE map display, flight strips, and D-BRITE provide a good picture of the current state of the airport. However, data regarding future departure demand on airport resources is not currently available.

NASA Ames Research Center, in cooperation with the Federal Aviation Administration (FAA) is developing a decision support tool known as the Surface Management System (SMS). The project is supported by NASA's Advanced Air Transportation Technologies (AATT) Project. SMS uses information provided by the new surface surveillance systems and departure plans provided by the air carriers in order to provide the Tower, TRACON, Center and air carriers with better information about current and future demand, thereby creating shared awareness of the departure situation and improving the capacity, efficiency, and flexibility of the airport¹.

SMS aids controllers with a variety of tasks including runway balancing and departure scenario optimization. Runway balancing is the task of ensuring that all active departure runways are equally busy in terms of imposed delay and usage. A departure scenario is defined as the mapping of departure fixes or gates[§] to a departure runway. The purpose of these runway assignment rules is to ensure that the airborne trajectories of aircraft that

*Aerospace Engineer.

[†]Human Factors Engineer.

[‡]Aerospace Engineer, Member AIAA.

[§] An aircraft's departure fix is the first fix listed in its flight plan. At DFW the 16 departure fixes are grouped into four "gates," one per side: North, South, East, and West.

takeoff from different runways do not cross. Departure scenario optimization is the task of ensuring that the current departure scenario provides the most efficient runway usage and leads to the least possible number of delays on the airport surface.

Occasionally, it is possible to further enhance the efficiency of the airport surface by identifying aircraft that should be exceptions to the departure scenario. In this case, SMS provides runway advisories via map displays and timelines, advising the controller to taxi the aircraft to an alternate departure runway. SMS also provides additional advisories to help manage surface movements and departure operations. For example, SMS aids controllers with the task of sequencing departure aircraft by taking into account inter-departure gaps required by wake-vortex considerations and downstream departure flow constraints.

SMS currently employs three types of user interfaces: map displays, timelines, and load graphs. Map displays of the airport surface provide a two-dimensional representation of the airport and include flight-specific information on data tags. Timelines provide flight-specific information and predictive time information, and load graphs provide aggregate data.

Two simulations were conducted in order to solicit user feedback about the SMS concept, the preliminary user interfaces, and the algorithm performance. These real-time controller-in-the-loop simulations of SMS were conducted in the Future Flight Central (FFC) air traffic control Tower simulation facility at NASA Ames Research Center in September, 2001 and January, 2002. FFC is a 360-degree, high fidelity control Tower simulator designed to provide the look and feel of a Level V airport Tower cab. Developed as a joint effort between NASA and the FAA, FFC uses twelve large rear projection screens and computer-generated imagery to provide a 360-degree out-the-window view. Controllers use standard headsets to talk to the pseudo-pilots who control the individual aircraft movements.

The initial simulation, held in September, 2001, consisted of three 45-minute runs. During each run, a different set of displays was presented to each Tower controller. Data were recorded during each run and the controllers completed questionnaires and participated in group debrief interviews after each run. Four active Dallas-Fort Worth (DFW) Tower controllers staffed the Local and Ground positions, while controllers from other airports observed and provided feedback. Observers were also present from Delta Air Lines and United Airlines. A large amount of expert user feedback was acquired through multiple discussions and informal debriefs with the controllers as well as

through questionnaires and recordings of the simulation proceedings.

The results of Simulation 1 indicated that map displays were well-liked by the Local and Ground controllers and that timelines had potential uses for them as well, but that both timelines and load graphs might be better suited for a Traffic Management Coordinator (TMC) or Supervisor position. The experimental design and results of Simulation 1 are described in detail in Reference 2. The feedback and human factors observations recorded during this simulation were incorporated into a refined version of SMS that was evaluated in a second simulation in January, 2002.

The second simulation consisted of nine hour-long runs based on actual DFW traffic. Data were recorded during each run and the controllers completed questionnaires after specific runs. Additionally, group debrief interviews were held after each run. More than 30 participants were involved in the simulation, including five active DFW Tower controllers, a controller from the Memphis, TN airport (MEM) Tower, a supervisor from the Norfolk, VA Tower, and airline observers from FedEx, Northwest Airlines, UPS, American Airlines, and United Airlines. The five active DFW controllers staffed the Tower positions, while the other controllers observed and provided feedback. A large amount of expert user feedback was acquired through the informal debrief sessions as well as through questionnaires and data collected during the simulation. This paper summarizes the methodologies employed, the SMS displays tested, and the human factors lessons learned from the simulation. SMS will be further refined based on the feedback from the simulation and will be evaluated next in the FedEx ramp tower at MEM in the summer of 2002. An evaluation in the ATC Tower will follow in 2003.

SIMULATION DESCRIPTION

Simulation Environment

In nominal operating conditions, DFW operates two air traffic control Towers, one controlling the west half of the airport, the other controlling the east half. Since the FFC facility can simulate only one Tower, only one side of the airport could be modeled. Therefore, since the majority of the gates and runways are located on the east half of the airport, a modified version of the East Tower of DFW was modeled. Operations were only conducted in South Flow under Visual Flight Rules (VFR) conditions. Figure 1 is a diagram of DFW airport. The box encloses the area that was modeled for the SMS simulation.




Figure 1: Layout of DFW

When DFW operates in South Flow, jet aircraft take off from runways 17R and 18L and land on runways 17C, 17L, and 18R. Runway 13L is used for prop departures and 13R is used for both prop and jet arrivals. As one of the goals of the simulation was to probe the efficacy of SMS in aiding the task of runway balancing, the procedures used on the east side of DFW needed to be altered to include two jet departure runways. Therefore, runway 17C was used for both arrivals and departures, thereby creating a second departure runway. Runway 13L was not used in the simulation.

Participants

Normal staffing of the DFW East Tower during busy traffic periods consists of eight positions: Supervisor, TMC, Clearance Delivery/Flight Data, Ground East 1 (Departure Ground or GE-1), Ground East 2 (Arrival Ground or GE-2), two Local controllers and a Cab Coordinator (CCE-1). Four of these eight positions, as well as an additional position created specifically for this simulation, were staffed by five active DFW controllers. The position that was created for the simulation was the “TMC Assistant” whose job it was to work with the TMC and to help evaluate the TMC displays. One of the active controllers, who is both a traffic manager and a controller at DFW, was assigned exclusively to the TMC position. The other four controllers rotated through the positions of Ground East 1 (GE-1), Ground East 2 (GE-2), Local East 1 (LE-1), and TMC Assistant. A retired controller staffed the Cab Coordinator (CCE-1) position throughout the simulation.

As a result of the altered DFW configuration, the roles of the Tower controllers changed slightly from their actual duties at DFW.

- During the simulation, as in current DFW operations, GE-1 was responsible for taxiing

departing flights from their spots¹ to their departure runway and sequencing the aircraft for departure. The difference between the role of GE-1 in day-to-day operations and his role during the simulation was that during the simulation this controller queued aircraft to depart off of 17C in addition to 17R and did not send any aircraft to the west side of the airport.

- At DFW, the GE-2 is responsible for taxiing arrivals that have cleared the active runways to their spot, mixing and merging flights coming to or leaving spot areas, and taxiing departures leaving from spots 50-53 to their departure runways. During the simulation, the GE-2 position incurred much less workload than the actual DFW position due to the fact that there were no aircraft arriving from the west side of the airport, as there are in current DFW operations.
- The LE-1 is responsible for all activity on runways 17R and 17C including providing aircraft with arrival and departure clearance and taxiing aircraft across active runways.
- CCE-1 is responsible for assisting LE-1 by scanning runways and taxiways for potential conflicts.
- The TMC monitors the departure scenario and runway balance to ensure that the airport is operating efficiently. The TMC is also responsible for the coordination among the controllers in the Tower. For the purpose of the simulation, the TMC was asked to determine the optimal time to switch operations on 17C from departures to arrivals and relay this information to the controllers.

Simulation Schedule

The simulation was comprised of nine hour-long runs. Three experimental conditions were tested and run three times each during the simulation. The traffic scenarios that were used approximated the 0800, 1130, and 1300 ‘rush’ periods at DFW. However, in order to minimize learning effects, the scenarios were presented as if they were all occurring in the 1130 time frame. Each rush was characterized by a departure push followed by an arrival rush, but the specific overlap times differed in each scenario.

The simulation began with an hour and a half of classroom training followed by one hour and 45 minutes of practice time in the FFC facility. When the

¹ A “spot” is the location on an airport surface at which an aircraft is transferred from ramp control to Tower control or vice versa.

training was complete, the nine simulation runs were conducted.

Experiment Design

SMS utilizes three types of displays to convey information and advisories: map displays, timelines, and load graphs. A map display is a two-dimensional representation of an airport outlining the airport surface (i.e., providing the outlines of the taxiways, runways, ramps) and showing the location and direction of travel for each aircraft.

Timelines, which are referenced to a physical location (e.g., a runway threshold or taxiway intersection), show the predicted times when aircraft will occupy that location but do not explicitly show the current location of each aircraft. Load graphs display the amount of present and forecasted demand on a specified airport resource (e.g., a runway, departure fix, or taxiway intersection). Load graphs display aggregate demand information rather than flight-specific information. Both timelines and load graphs have been used in the Center-TRACON Automation System (CTAS) Traffic Management Advisor (TMA) tool³. The SMS display types are described in greater detail in Reference 2.

Experimental Conditions

Three experimental conditions were tested during the simulation. Using feedback that was obtained during the initial simulation, specific displays were designed for each experimental condition and controller position. These conditions and the corresponding displays are described below.

Condition 1: Local and Ground displays

Condition 1, the Baseline Condition, was intended to replicate the workings of a control Tower without SMS. During this Condition, the Local and Ground controllers' displays consisted of an airport map display, similar to what they will have after ASDE-X is deployed in the Tower. The aircraft symbols on the map display were labeled with their flight IDs. Figure 2 is a portion of a sample map display. The controllers had the ability to toggle the flight IDs on or off independently for arrivals and departures to de-clutter the map. Controllers were also able to highlight the datablock for any aircraft by moving the mouse pointer over the aircraft symbol for that aircraft. Figure 3 is an example of a highlighted datablock.




Figure 2: Map Display



Figure 3: Highlighted Datablock

Condition 1: TMC displays

The TMC display during the Baseline Condition consisted of a load graph similar to the Enhanced Traffic Management System (ETMS) information that are currently provided in the Tower. The load graph shown in Figure 4 presented the upcoming arrival and departure demand in 15-minute intervals.

On the actual display, a white line (shown in the figure as a dark line) displays the predicted number of arrival aircraft, and a green line (shown in the figure as a dashed line) displays the predicted number of departure aircraft. The demand is presented in 15-minute periods, and the line represents the total number of aircraft that will be arriving or departing in that time period according to unconstrained demand information.




Figure 4: Sample ETMS-like load graph

Condition 2: Local and Ground displays

During Condition 2, the Local and Ground controllers were presented with airport map displays and timelines. The aircraft symbols on the map displays were labeled with expanded datablocks (i.e., flight specific

information such as aircraft type, departure gate and spot number was provided in datablocks). The specific information that was provided and the area covered by the map display depended on the controller position. Controllers had the ability to toggle through arrival and departure datablocks independently on the map display through three separate settings: no datablock, flight ID only, and expanded datablock. Figure 5 shows datablocks with flight-specific information. This datablock provides the flight ID and aircraft type on the first line and the initial departure fix and departure runway on the second line.



Figure 5: Datablock with Flight-Specific Information

The timeline displays that were shown to the Ground and Local Controllers differed according to controller position. The GE-1 had six timelines, four that were referenced to spot areas and two to the departure runways. The data tags on the timelines were color-coded according to departure gate. The GE-2 had four timelines that were referenced to the spot areas that were color-coded such that the arrival data tags were white and the departure data tags were green. The LE-1 had four timelines, one for each departure queue (EF, EG, EH, and the queue for runway 17C). Again, the data tags were color-coded by departure gate. All of the timelines for the Ground and Local Controllers had a ten minute look-ahead time.

Condition 2: TMC displays

During Condition 2, the TMC's display included timelines referenced to the runway thresholds, a map display covering the entire east side of the airport, and three load graphs: unconstrained arrival and departure demand, arrival and departure delay, and queue length to the departure runways.

Figure 6 is a sample of the arrival and departure delay load graph. On the actual display, a white line (shown in the figure as a dark line) represents the delay on the arrival aircraft, a green line (shown in the figure as a light line) represents the delay on the departure aircraft. A red horizontal line (not shown in the figure) represents a theoretical delay limit and can be placed by the controller.




Figure 6: Delay Load Graph provided to TMC

Figure 7 shows a sample arrival and departure demand load graph used by the TMC during Condition 2. On the actual display, a white line (shown in the figure as a dark line) represents the predicted arrival demand, and a green line (shown in the figure as a light line) represents the predicted departure demand.

In SMS, demand can be displayed as either constrained or unconstrained. Unconstrained demand means that the times that are being shown are calculated as if there were no other aircraft operating in the system. In other words, SMS does not run prediction algorithms to determine how the other aircraft will affect this flight's arrival or departure time. Constrained demand includes SMS predictions such as taxi delay and wake-vortex separations. The load graph in Figure 7 displays unconstrained demand.




Figure 7: Unconstrained Arrival and Departure Demand Load Graph provided to TMC

The TMC also used a load graph depicting the queue lengths at each of the runways. Figure 8 is a example of that departure queue load graph. The queue for 17R is depicted in yellow (shown in the figure as a dark line), and the queue for 17C is shown in blue (shown in the figure as a light line).




Figure 8: Departure Queue Load Graph

The TMC's timelines were referenced to the departure runways and displayed the expected arrival and departure traffic for those runways. Figure 9 is an example of the timelines that were presented to the TMC during Condition 2.

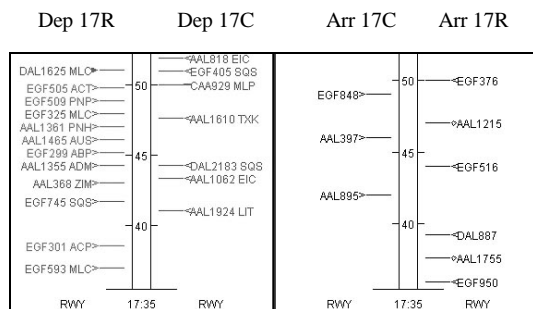


Figure 9: TMC Timelines

The TMC's timeline display provides future demand information for the three runways: 17R, 17C, and 17L. Departure information is displayed on the left timeline, and arrival information is displayed on the right. The left-hand side of the left timeline displays the departures from 17R, and the right-hand side of that same timeline displays the departures from 17C. On the right displays, the left-hand side displays the arrivals to 17C and the right-hand side displays the arrivals to 17L. The demand information that is provided on the timelines is constrained demand (i.e., it includes scheduling information provided by SMS).

Condition 3: Ground, Local, and TMC Displays

Condition 3 differed only slightly from Condition 2. All of the displays presented during Condition 2 were also used during Condition 3. The TMC received one additional display, the Configuration Change Advisory Tool.

The Configuration Change Advisory Tool, shown in Figure 10, provides the TMC with a suggested time to

switch runway 17C from a departure runway to an arrival runway. The tool shows the arrival, departure, and total delays as a function of the time at which the configuration is changed and selects the point at which the predicted total delay is lowest. In this example the tool has suggested 18:12 as the appropriate change time.




Figure 10: Sample Configuration Change Advisory Tool

Flight Strip Usage

During the week prior to the simulation, several of the DFW controllers participated in dry runs to assess the readiness of the simulation. During these dry runs, the DFW controllers expressed concern that working with both flight progress strips (FPS) and SMS created too much workload, as there were multiple sources for the same information. They were wary of spending too much time "heads-down" and asked if it would be possible to attempt running the simulation without the flight strips. After several trial runs without flight strips, the controllers requested that flight strips be removed when flight-specific information was presented on the map display, during Conditions 2 and 3.

It is important to note that it is not the intent of SMS to replace flight strips, and SMS does not replace several of the functions of flight strips. Air traffic controllers like flight strips for several reasons. The interface is familiar, easy-to-use, helps them instantly understand the current state of the traffic and lets them communicate without interrupting each other⁴.

The removal of flight strips during Conditions 2 and 3 introduced several issues. Controllers currently record clearances and departure queue assignments through various uses of flight strips, such as writing on them (see Fig. 11), "cocking" them, or sliding the strip out, each signifying something different. In this way, physical flight strips serve as a memory aid. Without the flight strips, the controllers were forced to find new

ways of recording clearances. Figure 11 is a flight strip from the SMS simulation. Several markings are visible on the strip. The 'Z' marking indicates that the aircraft was taxied on the Zulu taxiway, the '17C' marking indicates that the aircraft departed off of 17C, the 'DALL6' marking and slash through HUBB4 indicates that the aircraft's initial departure fix was changed from HUBB4 to DALL6, and the '-W' marking indicates that the Whiskey ATIS information was read to the pilot.

DAL1674	4331	DFW	DFW HUBB4-UTM ETC SHV		
T/B/722/G	P1815	2	DALL6	17C	-W
058	310				

Figure 11: Flight Strip from SMS Simulation

Data Collection

Several different types of data were recorded during the simulation. Visual and audio recordings were made of each run and SMS log files were recorded. SMS log files contain data such as the aircraft target positions, user keyboard entries, runway assignments and advisories, and runways used by each aircraft.

Human factors data were also collected via observation during each of the runs. Questionnaires were administered to controllers after the first and third run in each set of three runs. These questionnaires focused on the usability, suitability, and acceptability of the user interfaces. Usability refers to the ability of the controllers to readily obtain and use the information presented, suitability refers to the appropriateness of the user interfaces to the task requirements and information needs, and acceptability is the controller's trust in the information presented and his willingness to incorporate SMS into his/her task performance strategies. Most of the questions were based on a 7-point Likert scale. However, some multiple choice and open-ended questions were included. The questionnaires were designed to be specific to the experimental condition and controller position. Therefore, not all controllers were asked the same questions.

The usability, suitability, and acceptability questionnaires were administered to the Ground and Local controllers after each Baseline Condition and then again after the next two runs were completed (the next two runs consisted of Conditions 2 and 3, however not necessarily in that order). The TMC Assist position filled out questionnaires after each run, and the TMC filled out questionnaires after the first and last run in each experimental condition.

Controllers also provided feedback via informal 30-minute debrief sessions that were conducted after each run. During these debrief sessions, controllers were free to comment on any aspect of the simulation that

had just occurred, including aspects of the SMS user interface, traffic flow, etc. These discussions included requests for additional information, rationale for decisions made during each run, and requests for display changes. Finally, the TMC participated in structured interviews that focused on the decision process used to determine when to change the airport configuration and how SMS supports other TMC tasks.

HUMAN FACTORS RESULTS

The human factors work that was conducted during the simulation sought to better understand how controllers make decisions and how they interact with our automation and with each other. As a result, the lessons that were learned fall into two categories: controllers' opinions about the different display types and information about how controllers do their work, either in terms of thought processes or interactions.

For the purposes of this paper, only a subset of the data is described in detail. Although some numerical results are presented here, it is important to recognize that due to the small sample size these results represent trend information only.

Controller Opinions about Display Types

The controllers' responses to the displays were heavily dependent on the position that they were staffing. The tasks performed by the TMC are strategic in nature. The TMC is responsible for ensuring the smooth flow of traffic on the airport surface and for determining the appropriate configuration for the airport. Local and Ground controllers, on the other hand, perform more tactical tasks in order to direct the individual aircraft around the airport surface. As a result of the differences in their tasks, the displays that the TMCs preferred were different than the displays preferred by the Local and Ground controllers.

Ground and Local Responses: Map Displays

The feedback that was received about the map displays, both from the questionnaires and the informal debrief sessions, was positive. All of the Ground and Local controllers were in favor of having a map display that provided both aircraft location and flight-specific information (via datablocks). Additionally, the map display was the preferred information source for all tasks performed by GE-1, GE-2, LE-1 when it was presented with expanded datablocks. Controllers indicated that they trusted the information on the map display. When they were asked, "How much did you trust the information provided to you on the map display?," on a scale of 1="Trusted Completely" to 7="Did Not Trust at All," they responded with mean scores of:

GE-1: $x = 1.7, \sigma = 1.2$

GE-2: $x = 2.0, \sigma = 1.3$

LE-1: $x = 2.8, \sigma = 2.1$.

There are several hypotheses as to why controllers preferred using the map display to their other displays. First, the controllers who participated in the simulation were already familiar with map displays due to the fact that there is a map display currently in use the Tower at DFW. Also, map displays present an overhead view of the airport surface, much like the view the controllers see when they look out the windows of the Tower. Map displays present location information in a straightforward, intuitive manner. In addition, flight strips, which normally provide flight-specific information, were not available. The datablocks on the map display presented the easiest method of finding flight-specific information.

Several issues were identified with the map display, the most important of which was clutter. On many areas of the airport, such as departure queues, aircraft line up close to each other. When this occurs, the datablocks from the various aircraft overlap and become unreadable. Figure 12 is a sample of a map display presenting the departure queues to runway 17R.




Figure 12: Clutter on the Map Display

Several controllers offered suggestions for how to deal with the clutter issue. For example, the GE-2 suggested moving aircraft type to the second line of arrival datablocks (but leaving it on the first line of departure datablocks) to reduce clutter. One feature that was implemented prior to the simulation in order to combat the clutter issue was the ability to highlight an aircraft so that its datablock was more readable.

However, the controllers mentioned the highlighting feature was difficult to use. Instead they proposed a design in which relevant aircraft (i.e., the aircraft at the

front of the departure queues or the first aircraft waiting to cross an active runway) were always highlighted.

An issue that was mentioned by the GE-2 was that when the datablocks are set to display the flight ID only there is no cue as to which aircraft are arrivals and departures.

Ground and Local Responses: Timelines

The Ground and Local controllers' responses to the timeline displays were less favorable than those to the map displays. The overwhelming response was that they would like the information from the timelines to be presented on the map display. However this may be due to a need for a single display rather than due to the nature of timelines themselves. During the simulation, timelines were provided on a separate monitor from the map display, which provided the majority of the flight-specific information. The controllers were concerned about spending too much heads-down time. As a result, the timelines were used less frequently than the map display.

Local controllers found the timelines to be more useful than the Ground controllers did, as is evidenced by the results of the following question: "How useful was the information provided by the timelines?" on a scale of 1="Extremely Useful" to 7="Not at All Useful". The Ground controllers responded with mean responses of:

GE-1: $x = 5.7, \sigma = 0.6$

GE-2: $x = 5.0, \sigma = 2.0$

LE-1: $x = 3.3, \sigma = 1.5$.

This difference in ratings is directly related to the different tasks that are performed by the two different types of controllers. The Ground controllers used the timelines to inform them when an aircraft would be transitioning into or out of the active movement area. The Local controller used the bottom of the timelines to indicate which aircraft were available to be chosen next for departure. This information was also available on the local controller's map display, however due to the clutter issue described in the previous section, many datablocks were unreadable on the map display. Therefore, the timelines were valuable to the local controller who used them as a source for flight-specific information.

For the Ground controllers, the usefulness of timelines is dependent upon the amount of location information that they provide. The arrival timelines were less useful than departure timelines due to the lack of location information implicit in an aircraft's presence on the timeline. Whereas the departure timelines clearly indicated that the aircraft was moving in the ramp area and would soon be approaching the spot in order to

transition into the active movement area, an aircraft's presence on the arrival timelines only indicated that the aircraft would be approaching the spot sometime soon. The controller then had to search for the aircraft, which could have been located anywhere on the airport surface. The GE-2 suggested that timelines should be referenced to the crossing points as opposed to the spot areas in order to have implicit location information, as aircraft approaching the runway crossing points can only be taxiing on specific taxiways. The GE-2 liked using the timelines for reading the flight-specific information because of map clutter due to datablocks.

All of the Ground and Local controllers commented that the time duration of the timelines was unnecessarily large. Instead of having access to the predicted traffic over the next ten minutes, the Local controllers want to know the next one or two aircraft to depart or cross runways. As mentioned in the previous section, they suggested highlighting those aircraft on the map display instead of using timelines.

Another issue that was brought up by the controllers was the fact that aircraft that are located at gates far away from the spot appear earlier on the timelines because they push back earlier than aircraft that are parked close-by. Although the information provided on the timelines is accurate, and the aircraft is actually moving in the ramp area, it is not beneficial to the controller to be made aware of it when it is still several minutes away from crossing the spot.

The controllers also noted that the timelines were not sufficiently accurate. Timelines do not take into consideration aircraft that are not ready for pushback, gate arrival, etc, for various reasons. The GE-2 controller commented on some confusion because SMS provided inaccurate information when the aircraft taxied to a different parking gate than its assigned gate.

Ground and Local Responses: Runway Advisories

The concept of SMS providing runway assignment and departure sequence advisories was well received by the Ground controllers despite the fact that improvements to the current algorithms are necessary. The SMS algorithms that were used during the simulation did not consider the time needed for arrivals or departures to cross a runway, which led to advisories that did not reflect the way controllers would manage the aircraft. Despite these inaccuracies, the controllers stated that runway advisories freed up time for other tasks. Future research will focus on runway crossing delays.

The GE-1 controller responses indicated that runway advisories decreased their workload slightly and increased the efficiency of the airport surface slightly.

In response to the question "How much did using SMS runway advisories impact your workload?", the mean GE-1 response was $x = 2.7$, $\sigma = 1.5$ on a scale of 1="Decreased Workload" to 7="Increased Workload". Similarly, in response to the question "How much did the SMS runway advisories impact the efficiency of airport operations?" the mean GE-1 response was $x = 2.3$, $\sigma = 1.2$ on a scale of 1="Improved Airport Efficiency Greatly" to 7="Decreased Airport Efficiency Greatly". Additionally, in debrief sessions, the controllers said that SMS gave them a runway other than the one they would have preferred 50% of the time, nonetheless they claim to have followed the advisories 92% of the time.

In debriefing, the controllers stated that they would like the runway advisories to display a clear strategy. For example, during the simulation, the TMC made a decision to change runway 17C from a departure runway to an arrival runway at a specific time. The TMC advised the controllers that the switch time was approaching, and the controllers began to depart all subsequent departures off of 17R. However, the algorithm that was determining whether or not to present a runway advisory saw that there were still several minutes before 17C became an arrival runway and advised the controllers to take several more departures off of 17C, contrary to the instructions of the TMC. Therefore, according to the controllers, when the system makes the call to stop sending departures to runway 17C, it should be a distinct change. The controllers are under no obligation to follow the advisories, and after the flight enters the queue for another runway, the information in the datablock is updated to reflect the correct departure runway. This update only occurs after the flight has joined the other queue, and it may take several minutes for the information in the datablock to reflect the correct departure runway.

The Ground controllers also expressed that they liked the idea of having advisories, as long as they could override them whenever they wanted. Also, data show that controllers accepted more advisories during each successive run, which indicates increasing acceptance as their familiarity with advisories grew.

TMC Responses

The TMCs ranked timelines, Arrival/Departure Delay Load Graph, and Configuration Change Advisory Tool as the top three tools for performing their tasks. The map display was rarely used by the TMC.

The TMCs ranked the timelines as the most useful tool for determining the time at which to change the configuration of the airport. The timelines presented a

clear visual picture of the demand on each of the runways. The TMC used the timelines to see the departure demand and determine the appropriate time to switch the runway from a departures to arrivals. The TMCs also reported using the timelines tactically in order to find B757s or heavies to use to create gaps in the traffic flow to allow waiting aircraft to cross the runways. In this case, the TMC would identify an upcoming 757 and then inform the controller that the 757 should be used to cross multiple aircraft over the active runway.

The delay load graph was ranked as the second most important display because it supplements the timelines by providing information about how much delay is included into the predicted departure times. When a controller looks at a timeline of predicted traffic, it is impossible to determine how much delay each of the flights is absorbing. However, by correlating the demand shown on the timelines with the delay load graph, it is possible to determine if aircraft are being delayed or if they are scheduled to leave in the predicted traffic pattern.

The delay load graph and timelines were also used together to provide a “what-if” functionality that allowed the TMC to choose a switch time, evaluate the ramifications, and pick a new switch time to examine if necessary. The TMC commented that the delay load graph was the most helpful of all the load graphs and that the queue load graph was never used.

The third tool used by the TMC was the Configuration Change Advisory Tool, which informed the traffic manager of the theoretically optimal time to make the configuration switch. The TMC reported that the time was close, but not exactly the same as, the time that they would have chosen. The tool usually selected an earlier time than the TMC, which is logical due to the fact that the algorithms behind the configuration tool were not taking runway crossing times into account. It was, therefore, predicting lower taxi times overall and selecting an earlier configuration change time. The TMC noted that if the algorithms had taken runway crossing times into account, the tool would have been more accurate and, therefore, more useful.

Lessons about Controller Decision-Making

In addition to administering questionnaires and conducting interviews with the controllers, the human factors work conducted during the simulation involved observation (i.e., watching how the controllers worked and with whom they interacted). Although each controller has specific tasks that must be done individually, each controller pulls information from the others and from the various displays in the Tower.

Controllers collaborate with each other almost continuously. For instance, the Ground controllers, GE-1 and GE-2, spent much of their time observing the traffic on the LE-1 displays. Their goal was to monitor the LE-1’s traffic load and the demand on each of the departure fixes. They attempted to make decisions that would minimize workload on the next person downstream. The TMC also spent a large portion of time standing between the GE-1 and GE-2 positions, advising the controllers about tactical decisions. However, the LE-1 only seemed to interact with TMC on occasion and with GE-1 via receiving flight strips. The main implication that this has for SMS is that each of the displays needs to be clearly visible to each of the other controllers, and the designs need to be standardized in terms of color-coding and symbology.

The controllers told us that they do not follow any pattern of which aircraft go into which queue. However, in practice, they assign aircraft going to one fix to one queue and then put all “splitters[#]” on the other queues so that they can split up traffic as efficiently as possible. They view it as mixing up the traffic to operate as efficiently as possible. Different Ground controllers use different strategies to assign flights to departure queues. The goal is to set up the sequence so that no two flights in the final runway sequence have the same departure fix while also taking weight class into account.

Another observation that was made was that the controllers do not pay close attention to how long aircraft are waiting at spots or to cross runways. If an aircraft has been waiting a long time then the TMC usually points it out to them. Only if the aircraft queue began to back up did they seem to take any measures to cross a large number of aircraft.

SUMMARY

A major goal of the SMS Simulations was to learn about the roles of controllers, the methods they use to make decisions, and the types of information that they find useful. To this end, human factors studies investigated users’ preferences via questionnaires, interviews, and observational studies. Information was acquired about what displays may be appropriate for each user in a Tower environment. The roles of each controller were better defined, and the potential uses of SMS were explored.

It is clear that the Ground and Local controllers prefer using the map display for all of their tasks. The

[#] A splitter is an aircraft placed in a departure sequence to “split up” two other aircraft headed to the same departure fix.

controllers stated that timelines are not useful for them and should be removed from their displays, especially if a new function can be integrated such that relevant aircraft are highlighted on the map display. This would address several issues (including clutter) and would consolidate all pertinent information onto one monitor. However, much of the negative response to the timelines was influenced by the controllers' unfamiliarity with the displays and the timelines' inaccuracy due to the prototype SMS algorithms. It is important to note that until the SMS algorithms are refined and include estimates of runway crossing delays, timelines should not be dismissed as a potential display for Tower controllers. Future research will be conducted with accurate timelines in order to determine their potential uses.

Runway advisories were well-received by the Ground controllers. They appear to increase efficiency and decrease controller workload. However, the controllers stated that they must always be able to override the advisory and the advisories must display a clear strategy that is apparent to all users

It was determined that timelines are well suited to the strategic tasks of the TMC as are the Delay Load Graph and the Configuration Change Advisory Tool. However, demand load graphs and load graphs displaying departure queues are not useful to the TMC. It was also found that because the tasks conducted by the TMC are much more strategic in nature than those conducted by the Ground and Local controllers, a longer look-ahead time was appropriate.

Although each controller works individually, each of them also monitors the scenario downstream and takes the implications of their clearances on other controllers into consideration. Therefore, SMS displays must be consistent in terms of color-coding and symbology such that a controller glancing at another user's display is not confused by conflicting information.

Additionally, the use of flight strips and SMS together creates additional workload. More research will be conducted, either to determine how to design automation to support the use of strips without increasing workload or to create a system that can fully replace flight strips.

The human factors lessons learned from the second SMS simulation will be used to refine SMS. In addition, feedback from other user groups, such as ramp tower controllers, Airline Operations Center users, and TRACON and Center TMCs, will contribute to further development and refinement. SMS will first be demonstrated in the FedEx ramp tower at Memphis

airport in the summer of 2002. ATC Tower demonstrations will follow in 2003.

¹S. Atkins, C. Brinton, and D. Walton, "Functionalities, Displays, and Concept of Use for the Surface Management System," 21'st Digital Avionics Systems Conference, Irvine, CA, October 27-31, 2002.

²D. Walton and S. Atkins, "A Simulation Evaluation of the Surface Management System User Interfaces," Planned to be published as a NASA TM, 2002.

³K. Harwood and B. Sanford, "Denver TMA Assessment," NASA Contractor Report 4554, October 1993.

⁴W. MacKay, "Is Paper Safer? The Role of Paper Flight Strips in Air Traffic Control," ACM Transactions on Computer-Human Interaction, Vol. 6, No. 4, December 1999, pp 311-340.